

The hydrodynamic response of small-scale structure to reionization drives large IGM temperature fluctuations that persist to $z = 4$

Christopher Cain,¹★ Evan Scannapieco,¹ Matthew McQuinn,² Anson D’Aloisio³ and Hy Trac^{4,5}

¹School of Earth and Space exploration, Arizona State University, Tempe, AZ 85281, USA

²Department of Astronomy, University of Washington, Seattle, WA 98195-1580, USA

³Department of Physics and Astronomy, University of California, Riverside, CA 92521, USA

⁴McWilliams Center for Cosmology and Astrophysics, Department of Physics, Carnegie Mellon University, Pittsburgh, PA 15213, USA

⁵NSF AI Planning Institute for Physics of the Future, Carnegie Mellon University, Pittsburgh, PA 15213, USA

Accepted 2024 July 8. Received 2024 June 20; in original form 2024 May 3

ABSTRACT

The thermal history and structure of the intergalactic medium (IGM) at $z \geq 4$ is an important boundary condition for reionization, and a key input for studies using the Ly α forest to constrain the masses of alternative dark matter candidates. Most such inferences rely on simulations that lack the spatial resolution to fully resolve the hydrodynamic response of IGM filaments and minihaloes to H I reionization heating. In this letter, we use high-resolution hydrodynamic + radiative transfer simulations to study how these affect the IGM thermal structure. We find that the adiabatic heating and cooling driven by the expansion of initially cold gas filaments and minihaloes sources significant small-scale temperature fluctuations. These likely persist in much of the IGM until $z \leq 4$. Capturing this effect requires resolving the clumping scale of cold, pre-ionized gas, demanding spatial resolutions of $\leq 2 h^{-1}$ kpc. Pre-heating of the IGM by X-rays can slightly reduce the effect. Our preliminary estimate of the effect on the Ly α forest finds that, at $\log(k/[\text{km}^{-1}\text{s}]) = -1.0$, the Ly α forest flux power (at fixed mean flux) can increase ≈ 10 per cent going from 8 and $2 h^{-1}$ kpc resolution at $z = 4$ –5 for gas ionized at $z < 7$. These findings motivate more careful analyses of how the effects studied here affect the Ly α forest.

Key words: galaxies: high-redshift – radiative transfer – cosmology: reionization – dark matter.

1 INTRODUCTION

During the Epoch of Reionization, highly supersonic ionization fronts (I-fronts) increased the temperature of the intergalactic medium (IGM) by 1–3 orders of magnitude (Shapiro, Iliev & Raga 2004; Tittley & Meiksin 2007; Venkatesan & Benson 2011; D’Aloisio et al. 2019; Zeng & Hirata 2021). Measurements of the IGM temperature and its density dependence at $z \geq 4$ have begun to constrain the thermal history up to the Reionization epoch (Becker & Bolton 2013; Boera et al. 2019; Walther et al. 2019; Gaikwad et al. 2020; Wilson, Iršič & McQuinn 2022). These measurements constrain the reionization process itself (Nasir, Bolton & Becker 2016; Upton Sanderbeck, D’Aloisio & McQuinn 2016; Villasenor et al. 2022), which likely ended at $z \approx 5.5$ (Kulkarni et al. 2019; Keating et al. 2020; Nasir & D’Aloisio 2020; Bosman et al. 2022).

The main observable used to probe the IGM thermal history is the H I Ly α forest (e.g. Theuns et al. 2002). Ly α absorption is sensitive to the IGM gas temperature through both thermal broadening of the Ly α line and the temperature-dependent residual neutral fraction of ionized gas. In the low-density IGM (densities with respect to the cosmic mean of $\Delta \lesssim 10$) and well after reionization, the temperature–density relation (TDR) is often well-approximated in

reionization simulations (Kulkarni et al. 2015; Keating, Puchwein & Haehnelt 2018) by a power law of the form $T(\Delta) = T_0 \Delta^{\gamma-1}$, where T_0 is the temperature at mean density and $\gamma - 1$ is the power law index (Hui & Gnedin 1997; McQuinn & Upton Sanderbeck 2016). A number of studies have measured T_0 and γ using the $2 < z < 5$ Ly α forest (e.g. Becker et al. 2011; Boera et al. 2014; Hiss et al. 2018; Gaikwad et al. 2020).

Efforts to measure the thermal state of the IGM at $z > 4$ are important for several reasons. First, the IGM thermal history at these redshifts is a boundary condition for reionization, and can help distinguish between reionization scenarios with different timings and durations (Keating et al. 2020; Nasir & D’Aloisio 2020). Second, much of the constraining power for constraining alternative dark matter (DM) cosmologies using the Ly α forest comes from $z \approx 4$ –5 data (e.g. Viel et al. 2006, 2013; Baur et al. 2016; Iršič et al. 2017; Iršič, Xiao & McQuinn 2020). Such studies leverage the effects of the clumping properties of DM on small-scale structures in the Ly α forest (e.g. Iršič et al. 2024). These are sensitive to the IGM thermal history, so reliable inference requires reliable models for the thermal state of the gas.

A power-law TDR is motivated by heating/cooling processes that are important for low-density IGM gas. These include heating from photoionization and cooling from the expansion of the Universe and Compton scattering off the CMB (Hui & Gnedin 1997; McQuinn & Upton Sanderbeck 2016). However, one process that has received

* E-mail: clcain3@asu.edu

relatively little attention is the effect of pressure smoothing of the IGM by H_I reionization on the dynamics of the IGM thermal structure.¹ After I-fronts sweep through a region, gas structures at mass scales 10^4 – 10^8 M_⊙ undergo significant photoevaporation (Shapiro et al. 2004; Iliev, Scannapieco & Shapiro 2005; Ciardi et al. 2006; D’Aloisio et al. 2020; Nasir et al. 2021; Chan et al. 2024). This causes expansion of dense gas initially trapped in filaments and minihaloes, and compression in surrounding underdensities, which lead to adiabatic cooling and heating, respectively. Capturing these processes requires spatially resolving the Jeans scale in cold (≈ 10 – 1000 K), pre-ionized gas, which can be less than a ckpc (Gnedin 2000). Such resolution is hard to achieve in most Ly α forest simulations, which require volumes $\geq (20 \text{ cMpc})^3$ to capture the relevant large-scale fluctuations (Doughty et al. 2023). Indeed, Ly α forest convergence studies have only begun reaching the scales necessary to resolve some of this effect.² Some studies have found evidence for convergence in Ly α forest simulations (e.g. Bolton & Becker 2009), but often only in the limit that the IGM has been reionized for a long time, such that pressure smoothing takes place at redshifts much higher than are relevant for the Ly α forest.

Several recent studies have begun to converge on the small-scale dynamics of IGM gas in the aftermath of cosmological I-fronts. A number of works have focused primarily on characterizing the destruction of minihaloes and filamentary structures by reionization heating in the context of modelling the intergalactic ionizing opacity, and thus gave relatively little attention to the associated effects on temperature (e.g. Park et al. 2016; D’Aloisio et al. 2020; Nasir et al. 2021; Chan et al. 2024). The question of how this process affects the thermal structure of the IGM (in the context of the Ly α forest) was first addressed by Hirata (2018; see also Montero-Camacho, Zhang & Mao 2024). They found that pressure smoothing can induce significant and long-lived temperature fluctuations in the IGM, but only on the very small scales captured by their $\sim 300 h^{-1}$ kpc boxes. Here, we build on previous work in several crucial ways. First, we address the numerical convergence of the gas physics and its effects on the Ly α forest, which has yet to be studied down to the \sim kpc scales required to resolve pre-ionization clumping. Second, we use simulation volumes large enough ($L_{\text{box}} = 2 h^{-1}$ Mpc) to capture a representative sample of cosmic structures while simultaneously capturing much of the post-ionization pressure-smoothing effect. Finally, our simulations are performed with radiative transfer, which captures self-consistently the interplay between pressure smoothing, self-shielding, and the heating of IGM gas by I-fronts.

In this letter, we study the effect of pressure smoothing on the thermal state of the IGM down to $z = 4$ in different reionization scenarios, and comment on possible implications for the high-redshift Ly α forest. We use a suite of fully coupled hydrodynamic + radiative transfer (RT) simulations of IGM gas dynamics at $z = 4$ – 15 . These have sufficient spatial resolution to capture much of the response of minihaloes and filaments to the reionization process, and large enough volumes to at least capture some of the physical scales relevant for Ly α forest studies. We describe our numerical methods in Section 2, present our main results in Section 3, and conclude in Section 4. Throughout, we assume the following cosmological parameters: $\Omega_m = 0.305$, $\Omega_\Lambda = 1 - \Omega_m$, $\Omega_b = 0.048$, $h = 0.68$,

$n_s = 0.9667$ and $\sigma_8 = 0.82$, consistent with Planck Collaboration (2020) results. Distances are in co-moving units unless otherwise specified.

2 NUMERICAL METHODS

We simulated IGM gas dynamics using a modified version of the RADHYDRO code presented in Trac & Pen (2004) and Trac & Cen (2007). RADHYDRO solves the RT equation with ray tracing, which is fully coupled to the gas dynamics on a uniform Eulerian grid. DM is treated in the Lagrangian frame and gravity is calculated using a particle-mesh scheme. The simulations are initialized at $z = 300$ and run to $z = 4$. The code uses a reduced speed of light approximation to speed up the RT calculation, and it solves for the chemical and thermal evolution of the gas using a sub-cycling backwards-difference solver. All simulations are in $L_{\text{box}} = 2 h^{-1}$ Mpc boxes, and we have run simulations with $N = 1024^3$ DM particles, gas cells, and RT cells ($2 h^{-1}$ kpc cells). We also run $N = 256^3$ ($8 h^{-1}$ kpc cell) simulations to study resolution convergence. Our box size is chosen to get the largest possible statistical sample of cosmic structures whilst resolving most of the pressure smoothing effect. Our lower resolution is close to the $10 h^{-1}$ kpc resolution recommended by Doughty et al. (2023), which is on the high end of that achieved by most Ly α forest studies. So, they likely capture as much thermal structure in low-density gas as previous Ly α forest simulations. We refer to these as our ‘high-res’ and ‘low-res’ simulations, respectively.

Our simulation setup is similar to that of D’Aloisio et al. (2020), designed to model the dynamics of IGM gas after it is heated by external ionizing sources. We place $N_{\text{dom}} = 16^3$ RT domains on a regular grid and send plane-parallel rays from all six faces into each domain. Following D’Aloisio et al. (2020), the ionizing spectrum has a power law form, $J_\nu \propto \nu^{-\alpha}$ with $\alpha = 1.5$. The radiation turns on everywhere at a specified redshift z_{re} , with a flux density set to achieve a constant H_I photoionization rate in optically thin gas, $\Gamma_{-12} \equiv \Gamma_{\text{HI}}/10^{-12} \text{ s}^{-1}$. In this study, we will consider simulations with $\Gamma_{-12} = 0.3$, which is close to what is measured from the Ly α forest at $4 < z < 6$ (e.g. Becker & Bolton 2013; D’Aloisio et al. 2018; Bosman et al. 2022). During the few Myr the radiation takes to cross the RT domains, we ‘freeze’ hydrodynamics, gravity, and redshift evolution.³ We find that the gas temperature immediately following I-front passage (T_{reion}) in our simulations is in good agreement with the model of D’Aloisio et al. (2019; their equation 3). This guarantees that the hydrodynamic response of the gas to reionization is coherent everywhere in the box, with no unphysical gradients resulting from the RT domain structure. The methods and simulations will be detailed in a forthcoming paper.

To calculate Ly α forest statistics, we draw 100 random sightlines through the simulation volume that are long enough to cover the wavelength range of the Ly α forest (1025–1215 Å, or $\sim 0.5 h^{-1}$ Gpc). At $z = 4$, this is a total path length of $\approx 40 h^{-1}$ cGpc, sufficient to calculate converged Ly α forest statistics at scales captured by our boxes. Note that since the boxes are periodic, a single sightline wraps around the box many times. We include a condition on the angles of the random sightlines so that none of them point along a box axis, thus avoiding repeated structures.

¹Although see Puchwein et al. (2023) for a characterization of expanding filaments in the context of Ly α forest studies.

²To our knowledge, the highest resolution achieved in any Ly α forest convergence study (at the mean density) is $5 h^{-1}$ ckpc, in Doughty et al. (2023).

³Thermal evolution in cells less than 2 per cent neutral is also frozen during this time, which avoids unwanted cooling effects.

3 RESULTS

Fig. 1 visualizes the effect of pressure smoothing on the IGM density and temperature. The left and right halves show slices through the density and temperature at $z = 6$ and 4 (respectively) for low-res (top row) and high-res (middle row) simulations re-ionized at $z_{\text{re}} = 7$. The bottom panels show a $(200 \text{ } h^{-1}\text{kpc})^2$ zoom-in around an overdensity in the high-res simulation.

At $z = 6$ ($\Delta t \approx 170$ Myr after ionization), the IGM is responding hydrodynamically to reionization. The initially cold, dense filaments are being pressure smoothed by photoheating from reionization, driving compression of lower density gas surrounding them. The overlap of these expanding filaments generates an interference pattern in the density field that is conspicuous in the high-res simulations. The temperature maps show the resulting complex thermal structure. As filaments expand, their interiors are adiabatically cooled from $\approx 20\,000$ – $30\,000$ K to $\approx 5\,000$ – $10\,000$ K. On the other hand, the gas on the edge of the densest filaments is heated by compression and possibly weak shocks up to $\approx 50\,000$ K (Chan et al. 2024).⁴ This is much larger than the initial post I-front temperature of ~ 30000 K, and thus cannot be explained by radiative processes alone, which cool the gas at these temperatures. This process results in filaments having a ‘cored’ thermal structure (as noted by Ocvirk et al. 2016). Finally, in the low-density gas surrounding filaments, milder, but still significant, T fluctuations are driven by the expansion of smaller local overdensities ($\Delta \sim 1$ – 5). As a result, even the low-density IGM displays a complicated thermal structure. Note that our ‘freezing’ procedure, described in Section 2, guarantees that these effects are due to hydrodynamics and not gradients in the reionization redshift across the box (see also Appendix A for tests with flash-ionized runs).

The right panels show the same slice at $z = 4$ ($\Delta t \approx 780$ Myr). By this time, the expansion of filaments is nearly complete and the gas has cooled significantly. We see a stronger positive correlation between temperature and density than at $z = 6$. This is because the processes that set the power law TDR have had enough time to start dominating the thermal structure of the gas. However, pressure smoothing effects remain conspicuous in the temperature map. Compressed gas at the boundaries of overlapping filaments is hotter than average, and the cored thermal structure of filaments is still visible. This suggests that the imprint of pressure smoothing on the IGM thermal structure may persist well past the end of reionization. Visual comparison of the top and middle rows shows that these effects, though still visible in the low-res simulation, are significantly underresolved. This is because these simulations lack enough resolution to capture the *pre-ionized* sizes of the filaments, and thus miss much of the pressure smoothing caused by reionization.

Fig. 2 quantifies the evolution of the TDR. The top three rows show the $T - \Delta$ phase diagram at $z = 4, 5.5, 6$, and 6.9 (left to right) for simulations with $z_{\text{re}} = 7$. The 1st and 2nd rows show results for high and low-res simulations, respectively. The 3rd row shows high-res results for a simulation with a temperature floor of $T_{\text{min}} = 10^3$ K imposed at $z < 15$, well before z_{re} . This case is meant to roughly bracket possible effects of X-ray pre-heating on the pre-ionized structure of the IGM (Furlanetto, Oh & Briggs 2006; Fialkov,

Barkana & Visbal 2014).⁵ Such pre-heating can increase the Jeans scale of pre-ionized gas, resulting in less initial clumping and less pressure smoothing after I-fronts sweep through (D’Aloisio et al. 2020; Park et al. 2021). The 4th row shows the effect of T_{reion} on pressure smoothing. We show a simulation with a factor of 10 lower ionizing background of $\Gamma_{-12} = 0.03$ – in this run, the I-fronts moving much more slowly through the box initially, resulting in $T_{\text{reion}} \sim 20000$ K (at $\Delta = 1$). The 5th row shows a high-res simulation with $z_{\text{re}} = 9$ at $z = 4, 5.5, 8$, and 8.9 . The bottom row shows a high-res sim re-ionized extremely late ($z_{\text{re}} = 5$) at $z = 4.9, 4.7, 4.5$, and 4 . The black lines denote 1, 2, and 3σ contours of the distribution. We see that just after re-ionization (far right), the TDR is a tight, slightly inverted power law⁶. After $\Delta z = 1$, it evolves dramatically in the high-res runs. The expansion/compression processes drive a 1–1.5 dex scatter at 3σ in the temperature at mean density. The effect is much less significant in the low-resolution run, which displays only ≈ 0.5 dex scatter.

Comparing the 1st and 2nd rows confirms our earlier observation: that *the effects of hydrodynamics on the TDR are highly unconverged in our low-res simulations*. Since the initial sizes of filaments and minihaloes are not captured in these runs, their hydrodynamic response to reionization is not fully captured either. In fact, we have run a flash-ionized simulation with $1 \text{ } h^{-1}\text{kpc}$ resolution ($N = 2048^3$) with $z_{\text{re}} = 7$ down to $z = 6$ and found that even our high-res runs are still mildly unconverged in this process (at least, in the limit of no preheating – see Appendix A). This lack of convergence suggests that even $2 \text{ } h^{-1}\text{kpc}$ resolution may not be sufficient to fully capture the effect under study. Note that most of the scatter arises at densities $\Delta \lesssim 10$, which would not be star-forming gas – thus, the absence of star formation in our simulations should not affect this result. The third row shows that the effect is somewhat reduced if IGM pre-heating is significant. However, this case still displays much more thermal structure than the low-res run.⁷

The 4th row shows the effect of lowering T_{reion}^0 (that is, T_{reion} at $\Delta = 1$) from ≈ 28000 K to 20000 K. Based on Fig. 2 of D’Aloisio et al. (2019), these correspond to I-front speeds of $\sim 2 \times 10^4$ and 2×10^3 km/s for an $\alpha = 1.5$ spectrum, respectively, which roughly brackets the range expected during the latter half of reionization (see their fig. 7, and also Zeng & Hirata 2021). The TDR at $z = 6.9$ is inverted slightly more in this case. The spread in the TDR at $z = 6$ and 5.5 is slightly smaller than in the fiducial high-res case, but also slightly more than in the X-ray heated case. At $z = 4$, the scatter is similar to the X-ray case and slightly less than the fiducial case. Thus, we find that that differences in the I-front speed has an effect similar to (or less than) that of X-ray pre-heating.

The 5th row demonstrates that, given enough time, the TDR does eventually ‘relax’ to a power law with the expected slope. The $z_{\text{re}} = 9$ run displays the same behaviour at $z = 8$ that the $z_{\text{re}} = 7$ run does at $z = 6$, but by $z = 5.5$ the scatter is significantly lower, and at $z = 4$ ($\Delta t \approx 1$ Gyr from z_{re}), a tight power law has been achieved with a scatter of only ≈ 0.3 dex at 3σ . However in the $z_{\text{re}} = 7$ runs at $z = 4$ ($\Delta t \approx 800$ Myr), there is still significant scatter [≈ 0.8 (0.6)

⁴Through study of the entropy distribution of the gas, we find that most of the evolution immediately after reionization is consistent with adiabatic physics and not shocking, with only a small fraction of the gas showing significant entropy increases that evidence shocking.

⁵X-ray pre-heating likely would happen more quickly at lower redshifts compared to our thermal floor implementation, leaving less time for pressure smoothing before reionization and, hence, the gas would be clumpier at a given pre-reionization temperature.

⁶The inversion results from high-density gas cooling for longer inside I-fronts, reaching a lower post I-front temperature.

⁷We have also run simulations with $T_{\text{min}} = 100$ and 10 K. We find only a small difference for the former and no appreciable difference for the latter.




Figure 1. Visualization of the effects of pressure smoothing on the post-ionized IGM temperature. The left half shows density and temperature slices at $z = 6$ from a high-res simulation re-ionized at $z_{\text{re}} = 7$, and the right half shows the same for $z = 4$. The top row shows a low-res simulation (with $8 h^{-1} \text{ kpc}$ resolution), while the middle row shows a high-res ($2 h^{-1} \text{ kpc}$) simulation. The bottom row shows a $200 h^{-1} \text{ kpc}$ zoom-in around an overdensity in the high-res box. The dense gas inside the filaments is colder than its surroundings due to expansion cooling, while the surrounding gas has been heated by compression and (perhaps) weak shocks up to $\approx 50,000 \text{ K}$. Even the low-density gas between filaments displays significant temperature fluctuations due to the expansion of local overdensities. By $z = 4$ ($\Delta t \approx 780 \text{ Myr}$), the expanding filaments have overlapped and cooled significantly, and the temperature structure more closely resembles that of the density field. However, the effects of pressure smoothing (e.g. enhanced temperatures where expanding filaments overlap) remain conspicuous. Comparison of the top and middle rows shows that these effects are much less pronounced in the low-res simulation, suggesting a substantial lack of convergence in the IGM thermal structure at this resolution – a resolution similar to the highest resolution simulations of the Ly α forest. Indeed, noticeable differences persist to $z = 4$, suggesting that resolving these effects may be important for cosmological inference from the $z > 4$ Ly α forest.

dex for the high (low)-resolution case and ≈ 0.7 dex for the pre-heated model]. In the left-most row, the white dashed lines indicate the power law slope of the TDR at mean density expected from the analytical model of McQuinn & Upton Sanderbeck (2016). The $z_{\text{re}} = 9$ case closely follows this expectation, demonstrating that the heating/cooling processes responsible for it (e.g. photoheating and Compton cooling) have largely erased the temperature fluctuations caused by pressure smoothing. The bottom row ($z_{\text{re}} = 5$) contrasts the $z_{\text{re}} = 9$ scenario. In this case, the hydro-driven fluctuations in T peak at around $z = 4$. It is unclear how long such a patch would take to reach a tight power law. Patches of the IGM with $z_{\text{re}} < 6$ may fill up

to 20 per cent of the Universe in realistic late-reionization scenarios (Kulkarni et al. 2019; Keating et al. 2020; Nasir & D’Aloisio 2020).

These results may have complicating consequences for high-redshift Ly α forest studies that rely on concordance models of the IGM temperature. Most Ly α forest simulations upon which such studies are based have resolutions of $\Delta x_{\text{cell}} \geq 10 h^{-1} \text{ kpc}$ (at the mean density), which would likely cause them to miss much of the effect under study here. In principle, this could affect efforts to interpret IGM temperature measurements parametrized by a power law, and/or inferences based on the Ly α forest (e.g. Viel et al. 2013; Iršič et al.




Figure 2. Effect of hydrodynamic response to reionization on the TDR. Each panel shows the TDR in log colour-scale, with the black lines denoting the 1, 2, and 3σ contours. The top three rows show the TDR at several redshifts for $z_{\text{re}} = 7$. The first two rows compare the low and high-res simulations shown in Fig. 1, while the third row shows a high-res simulation in which we impose a minimum temperature floor of $T = 1000$ K at $z < 15$, well before z_{re} . This temperature floor mimics the possible smoothing effects of X-ray preheating, which may lessen the subsequent effect from reionization. The fourth row shows high-res results for $z_{\text{re}} = 9$. The bottom row shows the opposite extreme – a patch that re-ionizes very late at $z_{\text{re}} = 5$. The white lines in the left column show the predicted TDR slope from the analytical model of McQuinn & Upton Sanderbeck (2016). For $z_{\text{re}} = 9, 7$, and 5 , these slopes are $\gamma_{\text{pred}} = 1.57, 1.46$, and 1.18 , respectively. Note that the low-res case is characteristic of the best resolutions obtained in previous Ly α forest studies.

2024) that do not fully account for these effects. We emphasize that the effect of pressure smoothing on the IGM temperature structure is distinct from its effect on the density field itself, which has been well understood and characterized in the context of the Ly α forest (Oñorbe, Hennawi & Lukić 2017; Puchwein et al. 2023). It is also distinct from fluctuations on large scales caused by the well-studied patchy reionization effect (Trac, Cen & Loeb 2008; D’Aloisio, McQuinn & Trac 2015; Wu et al. 2019). Indeed, the degree of scatter in the TDR caused by the effect studied here is comparable to that found for patchy reionization by Puchwein et al. (2023, their fig. 5). However, it is unclear how these effects would interact since the pressure-smoothing effect has substantial z_{re} dependence.

Since our simulations are large enough to capture Ly α forest statistics at some of the scales relevant for such studies, we have made a preliminary effort to quantify the importance of these effects on the Ly α forest power spectrum at fixed mean flux.⁸ We find that in the extreme case with $z_{\text{re}} = 5$, the $z = 4$ Ly α flux power spectrum at $\log(k/[\text{km}^{-1}\text{s}]) = -1.0$ is ≈ 20 per cent higher in our high-res compared to our low-res simulations (with greater differences at larger k), suggesting significant lack of convergence. For $(z_{\text{re}}, z) = (7, 5)$, this discrepancy becomes ≈ 10 per cent. For $(z_{\text{re}}, z) = (9, 4)$

⁸We re-scale our Ly α opacities such that the mean transmission matches measurements from Becker & Bolton (2013) at $z = 4$ and Bosman et al. (2022) at $z = 5$.

and (7, 4), we find < 10 per cent disagreement between the low and high-res simulations. These differences become significantly larger, though, when comparing to idealized scenarios in which the temperature–density relation is assumed to be a tight power law. Note that the left column of Fig. 2 shows that for all cases except $z_{\text{re}} = 9$, the power-law parametrization represents the TDR poorly at $z = 4$. Ly α forest studies aimed at DM constraints typically use $-2.5 \lesssim \log(k/[\text{km}^{-1}\text{s}]) \lesssim -1.0$ in their analyses (Viel et al. 2013), indicated that these differences could be important. For Ly α forest analyses, these differences may be compensated by marginalization over thermal parameters during parameter inference. However, the thermal model that is used does not capture most of the dispersion in the TDR found in our high-resolution simulations and so it is not obvious that this marginalization is sufficient to robustly constrain cosmological parameters such as the DM mass. We plan to follow up with a more detailed, Ly α forest-focused study.

4 CONCLUSIONS

In this letter, we studied the effect of pressure smoothing from reionization on the IGM temperature–density relation. We found that the pressure smoothing of dense filaments and minihaloes and compression of voids caused by this process results in a complex thermal structure that can persist to at least $z = 4$ in much of the IGM. This structure differs substantially from the tight power law that is typically used to parametrize the thermal state of the low-density IGM in Ly α forest studies. This effect is somewhat reduced, but still considerable, in simulations that assume significant pre-heating by X-ray sources prior to reionization. We have demonstrated that simulations with $\geq 10 h^{-1}\text{kpc}$ spatial resolution (at mean density), upon which most Ly α forest inferences rely, miss much of this process because they do not resolve the initial sizes of cold gas structures prior to reionization. We have made a preliminary effort to quantify the effect on the Ly α forest flux power spectrum, estimating effects as large as several tens of per cent – with larger effects at higher k and lower z_{reion} .

Our preliminary results motivate several follow-up questions. First, it is unclear how these effects would interact with the well-studied patchy reionization effect (e.g. Trac et al. 2008; D’Aloisio et al. 2015; Wu et al. 2019; Puchwein et al. 2023). We have found that the effects of pressure smoothing on temperature are sensitive to the time at which the gas was re-ionized. It follows that this effect must be coupled to the large-scale patchiness of reionization. This patchiness might also affect the Ly α forest at much lower wavenumbers than studied here, though we cannot address this question here with our small, single-reionization-redshift boxes. Furthermore, our results might be affected by the inclusion of large-scale peculiar velocities, which are not well-captured in our small volumes and have been shown to be important for the Ly α forest (e.g. Molaro et al. 2022). It is also unclear whether these effects could bias cosmological inferences from the Ly α forest. This could be addressed in future work by combining small-box simulations at different z_{re} and box-scale densities (like the ones used in this work), and/or by achieving $\sim 2 h^{-1}\text{kpc}$ in $10\text{--}20 h^{-1}\text{Mpc}$ boxes.

ACKNOWLEDGEMENTS

We thank George Becker and Vid Iršič for helpful comments on the draft version of this manuscript. CC acknowledges support from the Beus Center for Cosmic Foundations. ES acknowledges support from NASA grant 80NSSC22K1265.

DATA AVAILABILITY

The data underlying this article will be shared upon reasonable request to the corresponding author.

REFERENCES

Almgren A. S., Bell J. B., Lijewski M. J., Lukić Z., Andel E. V., 2013, *ApJ*, 765, 39
 Baur J., Palanque-Delabrouille N., Yèche C., Magneville C., Viel M., 2016, *J. Cosmology Astropart. Phys.*, 2016, 012
 Becker G. D., Bolton J. S., 2013, *MNRAS*, 436, 1023
 Becker G. D., Bolton J. S., Haehnelt M. G., Sargent W. L. W., 2011, *MNRAS*, 410, 1096
 Boera E., Murphy M. T., Becker G. D., Bolton J. S., 2014, *MNRAS*, 441, 1916
 Boera E., Becker G. D., Bolton J. S., Nasir F., 2019, *ApJ*, 872, 101
 Bolton J. S., Becker G. D., 2009, *MNRAS*, 398, L26
 Bosman S. E. I. et al., 2022, *MNRAS*, 514, 55
 Chan T. K., Benítez-Llambay A., Theuns T., Frenk C., Bower R., 2024, *MNRAS*, 528, 1296
 Ciardi B., Scannapieco E., Stoehr F., Ferrara A., Iliev I. T., Shapiro P. R., 2006, *MNRAS*, 366, 689
 D’Aloisio A., McQuinn M., Trac H., 2015, *ApJ*, 813, L38
 D’Aloisio A., McQuinn M., Davies F. B., Furlanetto S. R., 2018, *MNRAS*, 473, 560
 D’Aloisio A., McQuinn M., Maupin O., Davies F. B., Trac H., Fuller S., Upton Sanderbeck P. R., 2019, *ApJ*, 874, 154
 D’Aloisio A., McQuinn M., Trac H., Cain C., Mesinger A., 2020, *ApJ*, 898, 149
 Doughty C. C., Hennawi J. F., Davies F. B., Lukić Z., Oñorbe J., 2023, *MNRAS*, 525, 3790
 Fialkov A., Barkana R., Visbal E., 2014, *Nature*, 506, 197
 Furlanetto S. R., Oh S. P., Briggs F. H., 2006, *Phys. Rep.*, 433, 181
 Gaikwad P. et al., 2020, *MNRAS*, 494, 5091
 Gnedin N. Y., 2000, *ApJ*, 542, 535
 Hirata C. M., 2018, *MNRAS*, 474, 2173
 Hiss H., Walther M., Hennawi J. F., Oñorbe J., O’Meara J. M., Rorai A., Lukić Z., 2018, *ApJ*, 865, 42
 Hui L., Gnedin N. Y., 1997, *MNRAS*, 292, 27
 Iliev I. T., Scannapieco E., Shapiro P. R., 2005, *ApJ*, 624, 491
 Iršič V. et al., 2017, *Phys. Rev. D*, 96, 023522
 Iršič V., Xiao H., McQuinn M., 2020, *Phys. Rev. D*, 101, 123518
 Iršič V. et al., 2024, *Phys. Rev. D*, 109, 043511
 Keating L. C., Puchwein E., Haehnelt M. G., 2018, *MNRAS*, 477, 5501
 Keating L. C., Weinberger L. H., Kulkarni G., Haehnelt M. G., Chardin J., Aubert D., 2020, *MNRAS*, 491, 1736
 Kulkarni G., Hennawi J. F., Oñorbe J., Rorai A., Springel V., 2015, *ApJ*, 812, 30
 Kulkarni G., Keating L. C., Haehnelt M. G., Bosman S. E. I., Puchwein E., Chardin J., Aubert D., 2019, *MNRAS*, 485, L24
 McQuinn M., Upton Sanderbeck P. R., 2016, *MNRAS*, 456, 47
 Molaro M. et al., 2022, *MNRAS*, 509, 6119
 Montero-Camacho P., Zhang Y., Mao Y., 2024, *MNRAS*, 529, 3666
 Nasir F., D’Aloisio A., 2020, *MNRAS*, 494, 3080
 Nasir F., Bolton J. S., Becker G. D., 2016, *MNRAS*, 463, 2335
 Nasir F., Cain C., D’Aloisio A., Gangolli N., McQuinn M., 2021, *ApJ*, 923, 161
 Oñorbe J., Hennawi J. F., Lukić Z., 2017, *ApJ*, 837, 106
 Ocvirk P. et al., 2016, *MNRAS*, 463, 1462
 Park H., Shapiro P. R., Choi J.-h., Yoshida N., Hirano S., Ahn K., 2016, *ApJ*, 831, 86
 Park H., Shapiro P. R., Ahn K., Yoshida N., Hirano S., 2021, *ApJ*, 908, 96
 Planck Collaboration, 2020, *A&A*, 641, A6
 Puchwein E. et al., 2023, *MNRAS*, 519, 6162
 Shapiro P. R., Iliev I. T., Raga A. C., 2004, *MNRAS*, 348, 753
 Theuns T., Zaroubi S., Kim T.-S., Tzanavaris P., Carswell R. F., 2002, *MNRAS*, 332, 367

Tittley E. R., Meiksin A., 2007, *MNRAS*, 380, 1369
 Trac H., Cen R., 2007, *ApJ*, 671, 1
 Trac H., Pen U.-L., 2004, *New Astron.*, 9, 443
 Trac H., Cen R., Loeb A., 2008, *ApJ*, 689, L81
 Upton Sanderbeck P. R., D'Aloisio A., McQuinn M. J., 2016, *MNRAS*, 460, 1885
 Venkatesan A., Benson A., 2011, *MNRAS*, 417, 2264
 Viel M., Lesgourgues J., Haehnelt M. G., Matarrese S., Riotto A., 2006, *Phys. Rev. Lett.*, 97, 071301
 Viel M., Becker G. D., Bolton J. S., Haehnelt M. G., 2013, *Phys. Rev. D*, 88, 043502
 Villasenor B., Robertson B., Madau P., Schneider E., 2022, *ApJ*, 933, 59
 Walther M., Oñorbe J., Hennawi J. F., Lukic Z., 2019, *ApJ*, 872, 13
 Wilson B., Iršič V., McQuinn M., 2022, *MNRAS*, 509, 2423
 Wu X., McQuinn M., Kannan R., D'Aloisio A., Bird S., Marinacci F., Davé R., Hernquist L., 2019, *MNRAS*, 490, 3177
 Zeng C., Hirata C. M., 2021, *ApJ*, 906, 124

APPENDIX A: COMPARISON TO NYX & RESOLUTION CONVERGENCE

In this appendix, we compare our simulations with similar runs carried out with the Nyx cosmological hydrodynamic code (Almgren et al. 2013) and assess resolution convergence. We have run Nyx simulations using the same initial conditions in our RADHYDRO runs with $2 h^{-1}$ kpc resolution, with flash re-ionization at $z = 7$. In Nyx, we can set the heat injection from reionization to produce a constant T_{reion} . We find that $T_{\text{reion}} = 32,000$ K reproduces the initial temperature of low-density gas in RADHYDRO reasonably well for $z_{\text{re}} = 7$. Note that this setup neglects the slightly lower T_{reion} values in filaments.

The top two rows of Fig. A1 show the TDR at $z = 5$ and 6 for RADHYDRO and Nyx at our fiducial $2 h^{-1}$ kpc resolution. In the Nyx panels, the cyan lines show the corresponding RADHYDRO contours to aid the eye in comparison. We find good agreement in the TDR between the two codes. One difference is that the T distribution close to the mean density is slightly wider in RADHYDRO, with more gas getting below 10^4 K. This may be because the dense filaments start out colder in RADHYDRO than Nyx, since RADHYDRO accounts for the density dependence of T_{reion} (owing to its full RT treatment). The densest structures ($\Delta \gtrsim 100$) are also able to self-shield against ionizing radiation for 10s to 100s of Myr after I-fronts pass through, and thus can remain cold long after the box ionizes.

The bottom row shows the effect of spatial resolution. The left and right panels show the TDR for our low-res simulation and that of a flash-ionized (no RT) run with $1 h^{-1}$ kpc resolution. The cyan curves denote the $2 h^{-1}$ kpc-resolution contours. As seen in the main text, our low-res runs are very unconverged in the TDR. The $1 h^{-1}$ kpc simulation displays a slightly wider spread of temperatures, particularly close to the mean density. This suggests that even our

high-res runs may not be fully converged in the behaviour of the TDR on small scales. This is unsurprising, since the characteristic sizes of pre-ionized structures can be $1 h^{-1}$ kpc or smaller if the gas is sufficiently cold. However, we note that convergence criteria are likely to be most strict in these runs without X-Ray pre-heating as the gas near the mean density has cooled to ~ 1 K by these redshifts (resulting in it being maximally clumpy). Including the (highly uncertain) effects of pre-heating may eliminate the smallest gas structures and ease convergence criteria. Thus, our convergence analysis reflects an upper limit on the resolution requirements of the TDR.




Figure A1. Comparison of the TDR in RADHYDRO and Nyx (top and middle row). We flash-ionize a Nyx box at $z = 7$ and with $T_{\text{reion}} = 32,000$ K to codes at $z = 5$ and 6 with $2 h^{-1}$ kpc resolution. In the Nyx panels, the cyan lines denote the corresponding RADHYDRO contours. The agreement is excellent overall at both redshifts, the only notable difference being that RADHYDRO has a slightly wider distribution of temperatures near the mean density. This is explainable by differences in the T_{reion} physics, as explained in the text. The bottom row shows our low-res simulation (left) and our flash-ionized box with $1 h^{-1}$ kpc resolution (right) at $z = 6$, with the cyan contours denoting our $2 h^{-1}$ kpc run. As seen in the main text, the former is dramatically unconverged. The latter displays a slightly wider spread of temperatures than the $2 h^{-1}$ kpc case (cyan contours), particularly at densities slightly close to the mean. This suggests that even our high-res runs may not be fully converged (at least, in the limit of no pre-heating).

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.